Options For This Howto

- Free Support
- Paid Support

Navigation

- Howtos
  - Linux
    - CentOS
    - Debian
    - Fedora
    - Kernel
    - Mandriva
    - PCLinuxOS
    - SuSE
    - Ubuntu
  - Apache
  - Backup
  - Control Panels
  - DNS
  - Desktop
  - Email
    - Anti-Spam/Virus
    - Postfix
    - Procmail
    - Sendmail
  - FTP
  - Monitoring
  - MySQL
  - Programming
    - C/C++
    - PHP
  - Samba
  - Security
    - Anti-Spam/Virus
  - Virtualization
  - Other
  - FreeBSD
  - Commercial
- Mini-Howtos
- Forums
- Screenshots
- News
- Search
- Contribute
- Membership
- Login
- Site Map
- Advertise
- Contact
- Disclaimer
- Imprint

Username:

Password:

Remember Me?

Log in

Create a new account
Request new password

Screenshots

- Asianux 2.0
- Blag Linux 30001
- Blag Linux 49999 (alpha)
- CentOS 4.2
- Debian Sid (Xfce)
- Fedora Core 4
- Fedora Core 5
- Fedora Core 6
- IPCop 1.4.10
- Kanotix 2005-03
- Knoppix 4.02
- Kubuntu 5.10 RC1
- Kurumin Linux 6.0
- Mandriva 2006 Powerpack+
- Mandriva 2006 RC2
- PCLinux OS 0.92
- Slackware 10.2
- SuSE 10.0 (OSS)
- Symphony OS BETA1 PR1
- Ubuntu 5.10 Preview
- Ubuntu 6.06 LTS
- Xandros 3.0.2 Open Circulation Edition

Newsletter

Subscribe to HowtoForge Newsletter and stay informed about our latest HOWTOs and projects.

enter email address

Submit

(To unsubscribe from our newsletter send a blank email with the subject 'unsubscribe' to: newsletter @howtoforge.com.)

Who's online

There are currently 13 users and 2075 guests online.

HowtoForge Forums

Cant Find welcome mail

Mailuser and Aliases

internal dummy connection

mod_security breaking ispc

How to change file rights?

Viewing awstats

Proftpd-Mysql on Ubuntu Breezy

ISP Donfig IP setting question.

PHP binary does not work

Changin to Shell access for a users does'nt ...

Browse archives

« August 2007

Mo	Tu	We	Th	Fr	Sa	Su
		1	2	3	4	5
6	7	8	9	10	11	12
13	14	15	16	17	18	19
20	21	22	23	24	25	26
27	28	29	30	31		

News

Interview: Matthew Dillon

Second Life coders challenge to rebuild Sentenal

Migrating from Windows to Linux

Press Release: Lenovo and Novell Unveil Linux-enabled ThinkPad Mobile Workstation

Press Release: Lenovo and Novell Unveil Linux-enabled ThinkPad Mobile Workstation

Google Using OpenMoko For gPhone?

Monitor Your Systems With FreeIPMI

FreeBSD Apply Binary Security Updates and Patch management

Granny Uses PCLinuxOS

Lenovo to offer SUSE Linux Preload on ThinkPad Notebooks

more

Recent comments

command missing after mysql password set
6 hours 31 min ago

mysqldump use
15 hours 16 min ago

Thanks for your commitment!
15 hours 45 min ago

Nice article
1 day 19 hours ago

I can't believe
3 days 2 hours ago

Howto?
3 days 19 hours ago

If you are on Debian or
4 days 16 hours ago

On Ubuntu Feisty
4 days 20 hours ago

great as far as it goes...
5 days 3 hours ago

Non supported Ati cards
5 days 6 hours ago

Syndicate

XML

You are here: **Home** » **Howtos** » **MySQL** » **How To Set Up A Load-Balanced MySQL Cluster**

How To Set Up A Load-Balanced MySQL Cluster

Submitted by falko on Mon, 2006-03-27 11:00. :: MySQL

How To Set Up A Load-Balanced MySQL Cluster

Version 1.0
Author: Falko Timme <ft [at] falkotimme [dot] com>
Last edited 03/27/2006

This tutorial shows how to configure a MySQL 5 cluster with three nodes: two storage nodes and one management node. This cluster is load-balanced by a high-availability load balancer that in fact has two nodes that use the `Ultra Monkey` package which provides `heartbeat` (for checking if the other node is still alive) and `ldirectord` (to split up the requests to the nodes of the MySQL cluster).

In this document I use Debian Sarge for all nodes. Therefore the setup might differ a bit for other

distributions. The MySQL version I use in this setup is 5.0.19. If you do not want to use MySQL 5, you can use MySQL 4.1 as well, although I haven't tested it.

This howto is meant as a practical guide; it does not cover the theoretical backgrounds. They are treated in a lot of other documents in the web.

This document comes without warranty of any kind! I want to say that this is not the only way of setting up such a system. There are many ways of achieving this goal but this is the way I take. I do not issue any guarantee that this will work for you!

1 My Servers

I use the following Debian servers that are all in the same network (192.168.0.x in this example):

- *sql1.example.com*: *192.168.0.101* MySQL cluster node 1
- *sql2.example.com*: *192.168.0.102* MySQL cluster node 2
- *loadb1.example.com*: *192.168.0.103* Load Balancer 1 / MySQL cluster management server
- *loadb2.example.com*: *192.168.0.104* Load Balancer 2

In addition to that we need a virtual IP address : *192.168.0.105*. It will be assigned to the MySQL cluster by the load balancer so that applications have a single IP address to access the cluster.

Although we want to have two MySQL cluster nodes in our MySQL cluster, we still need a third node, the MySQL cluster management server, for mainly one reason: if one of the two MySQL cluster nodes fails, and the management server is not running, then the data on the two cluster nodes will become inconsistent ("*split brain*"). We also need it for configuring the MySQL cluster.

So normally we would need five machines for our setup:

2 MySQL cluster nodes + 1 cluster management server + 2 Load Balancers = 5

As the MySQL cluster management server does not use many resources, and the system would just sit there doing nothing, we can put our first load balancer on the same machine, which saves us one machine, so we end up with four machines.

2 Set Up The MySQL Cluster Management Server

First we have to download MySQL 5.0.19 (the **max** version!) and install the cluster management server (*ndb_mgmd*) and the cluster management client (*ndb_mgm* - it can be used to monitor what's going on in the cluster). The following steps are carried out on *loadb1.example.com* (*192.168.0.103*):

loadb1.example.com:

```
mkdir /usr/src/mysql-mgm
cd /usr/src/mysql-mgm
wget http://dev.mysql.com/get/Downloads/MySQL-5.0/mysql-max-5.0.19-linux-i686-\
glibc23.tar.gz/from/http://www.mirrorservice.org/sites/ftp.mysql.com/
tar xvfz mysql-max-5.0.19-linux-i686-glibc23.tar.gz
cd mysql-max-5.0.19-linux-i686-glibc23
mv bin/ndb_mgm /usr/bin
mv bin/ndb_mgmd /usr/bin
chmod 755 /usr/bin/ndb_mg*
cd /usr/src
rm -rf /usr/src/mysql-mgm
```

Next, we must create the cluster configuration file, */var/lib/mysql-cluster/config.ini*:

loadb1.example.com:

```
mkdir /var/lib/mysql-cluster
cd /var/lib/mysql-cluster
vi config.ini
```

```
[NDBD DEFAULT]
NoOfReplicas=2

[MYSQLD DEFAULT]

[NDB_MGMD DEFAULT]

[TCP DEFAULT]

# Section for the cluster management node
[NDB_MGMD]
# IP address of the management node (this system)
HostName=192.168.0.103
```

Please replace the IP addresses in the file appropriately.

Then we start the cluster management server:

loadb1.example.com:

```
ndb_mgmd -f /var/lib/mysql-cluster/config.ini
```

It makes sense to automatically start the management server at system boot time, so we create a very simple init script and the appropriate startup links:

loadb1.example.com:

```
echo 'ndb_mgmd -f /var/lib/mysql-cluster/config.ini' > /etc/init.d/ndb_mgmd
chmod 755 /etc/init.d/ndb_mgmd
update-rc.d ndb_mgmd defaults
```


- How To Set Up A Load-Balanced MySQL Cluster - Page 2
- How To Set Up A Load-Balanced MySQL Cluster - Page 3
- How To Set Up A Load-Balanced MySQL Cluster - Page 4
- How To Set Up A Load-Balanced MySQL Cluster - Page 5
- How To Set Up A Load-Balanced MySQL Cluster - Page 6
- How To Set Up A Load-Balanced MySQL Cluster - Page 7
- How To Set Up A Load-Balanced MySQL Cluster - Page 8

next
How To Set Up A Load-Balanced MySQL Cluster
- Page 2

Copyright © 2006 Falko Timme
All Rights Reserved.

print: this | all page(s) | login or register to post comments | email this page | view as pdf |

Related Tutorials

- How To Set Up Database Replication In MySQL
- How To Set Up A Loadbalanced High-Availability Apache Cluster
- Setting Up A Highly Available NFS Server

Comment viewing options

Threaded list - expanded | Date - newest first | 10 comments per page | Save settings

Note: Select your preferred way to display the comments and click "Save settings" to activate your changes.

⚠ *Please do not use the comment function to ask for help! If you need help, please use our forum: http://www.howtoforge.com/forums*

⚠ *Comments will be published after administrator approval.*

multiple management servers

Submitted by broadcastwebs on Thu, 2007-05-31 15:28.

You can make each balancer a management server and eliminate a single point of failure.

Install manager on both lb and add both to config.ini

```
# Managment Server 1
[NDB_MGMD]
HostName=192.168.0.8                # the IP of the First Management Server
ID=1
Datadir=/var/lib/mysql-cluster

# Managment Server 2
[NDB_MGMD]
HostName=192.168.0.9                # the IP of the Second Management Server
ID=2
Datadir=/var/lib/mysql-cluster
```

then on each data node modify my.cnf

```
[mysqld]
ndbcluster
ndb-connectstring = "host=192.168.0.8,host=192.168.0.9"

[ndb_mgm]
connect-string = "host=192.168.0.8,host=192.168.0.9"

[ndbd]
connect-string = "host=192.168.0.8,host=192.168.0.9"
```

make sure to run ndbd --initial

login or register to post comments | email this page | view as pdf

src package in /usr

Submitted by Anonymous on Sun, 2006-04-23 16:33.

You use Debian and install package from source. Don't store files in /usr/bin. Use /usr/local or /opt.

Custom packages installed in /usr can be broken by Debian packaging system.

login or register to post comments | email this page | view as pdf

Actually...

Submitted by Anonymous on Sun, 2006-08-27 01:30.

Actually, if you use chkinstall instead of 'make instal', it'll add the package to your apt setup so the files won't be overwritten

login or register to post comments | email this page | view as pdf

Single point of failure?

Submitted by Anonymous on Wed, 2006-03-29 19:44.

The cluster management software seems to be a single point of failure; that is, if the load balancer running this software goes down, doesn't the cluster either go down or end up with inconsistent data ("*split brain*", as referenced in the article)?

I'm very new to clustering, so I'd be happy to learn why I'm wrong!

login or register to post comments | email this page | view as pdf

No ...

Submitted by Anonymous on Fri, 2006-05-05 13:24.

The storage and MySQL Server nodes are not dependent on the management server for their execution. Its purpose is only to manage the cluster. It may fail and be restarted any number of times without affecting the running MySQL Cluster.

login or register to post comments | email this page | view as pdf

Memory Only Storage

Submitted by Anonymous on Wed, 2006-03-29 00:51.

Memory only storage is a significant limitation. I hope this is fixed in a future version.

login or register to post comments | email this page | view as pdf

Memory Only Storage

Submitted by Anonymous on Sun, 2006-04-23 17:13.

"Memory only storage is a significant limitation. I hope this is fixed in a future version." This isn't a limitation to be

fixed, but the fundamental tradeoff in MySQL Cluster Server's design: by accepting the limitation of being memory-based instead of disk based, it can be several orders of magnitude faster. If your data can't fit in RAM, and you don't need the performance, you should use one of the disk-based table types.

login or register to post comments | email this page | view as pdf

Not anymore

Submitted by Anonymous on Sun, 2006-04-23 16:32.

Well, not anymore:

In MySQL 5.1, the memory-only requirement of MySQL Cluster is removed and operational data may now be accessed both on disk and memory. A DBA can specify that table data can reside on disk, in memory, or a combination of main memory and disk (although a single table can only be assigned to either disk or main memory). Disk-based support includes new storage structures - tablespaces – that are used to logically house table data on disk. In addition, new memory caches are in place to manage the transfer of data stored in tablespaces to memory for fast access to repeatedly referenced information.

login or register to post comments | email this page | view as pdf

MySQL cluster - no production ready

Submitted by Anonymous on Mon, 2006-03-27 19:10.

This is rather unfortunate, but without foreign keys support and memory-only storage MySQL cluster is not a really viable solution for most RDBMS users.

login or register to post comments | email this page | view as pdf

foreign keys support

Submitted by Anonymous on Tue, 2006-03-28 23:06.

The InnoDB engine supports foriegn keys and works with the MySQL cluster so your comment is incorrect sir.

login or register to post comments | email this page | view as pdf

1 2 next page last page

Copyright © 2007 HowtoForge - Linux Howtos and Tutorials
All Rights Reserved.